

REMARKS

Applicant thanks Examiner Ly for a helpful telephone interview with the undersigned and Dr. Dally on December 19, 2003. In that interview, the previously filed response was discussed and it was agreed that the claims were allowable over the then cited references of Ben-Michael and Brant, but the Examiner indicated that he would conduct a further search.

The drawings were objected to for failure to include the legend "prior art" in Fig. 4. Reconsideration of that requirement is requested. Although Fig. 4 refers to a router developed by the Assignee of this application prior to the present invention, it is not prior art. The router presented in Fig. 4 was not introduced in a product prior to May 21, 1998, that is, one year prior to the filing of the present invention.

Claims 1-4, 8-19, 23-34 and 36-42 and 44-53 were rejected under 35 U.S.C. 103(a) as being unpatentable over Kadambi et al. (hereinafter Kadambi), Ben-Michael et al. (Ben-Michael) and further in view of Brant et al. (Brant). That rejection is respectfully traversed and reconsideration is requested.

The Kadambi application describes an Ethernet switch that supports both 10/100Mbit Ethernet (via EPICs 20a, 20b, 20c) and Gbit Ethernet (via GPICs 30a, 30b). Arriving Ethernet packets are buffered either in on-chip memory (CBP 50) or off-chip memory (GBP 60). Both of these memories are global in the sense that they are shared across all of the ports (EPICs and GPICs). (By contrast, see new claims 54-56.) For each packet, the algorithm of Figure 12, described in paragraphs 0116, 0117, and 0217 to 0220, determines whether the packet is buffered to on-chip memory (CBP) or off-chip memory (GBP) by comparing the packet length and the cell count for the ingress channel against a set of pre-determined limits. Unlike the claimed invention, the system of the Kadambi application does not evict packets from on-chip to off-chip memory; the on-chip memory does not function as a cache. Once a packet is stored in the on-chip memory (CBP), it remains there until it is transmitted. When a packet from off-chip memory (GBP) is to be transmitted, it is first read into on-chip memory (CBP) (paragraph 0208).

-12-

The Examiner recognizes that Kadambi does not disclose the information units from the first set of buffers being evicted to a second set of buffers. For that feature, Ben-Michael has been cited. However, it is respectfully submitted that there would be no reason for one skilled in the art to combine the references. The Kadambi reference stores a packet either to on-chip memory or to off-chip memory when it arrives and never evicts on-chip packets to off-chip memory. The Ben-Michael reference, when using off-chip memory, transfers every token from on-chip memory (the tail of the queue) to off-chip memory (the middle of the queue) and back to on-chip memory (the head of the queue). These transfers are all made in FIFO order. Each of the references describes a different method for managing storage, and the two methods are mutually incompatible.

Further, if one were to combine the two references as suggested by the Examiner, one would make use of the FIFO approach of Ben-Michael in the Kadambi reference, still failing to provide a "cache with information units being evicted to the second set of buffers according to an algorithm other than order of receipt in the first buffer." For this latter feature, the Examiner has cited yet a third reference, Brant. However, as successfully argued in the last response, there is no suggestion in the prior art of combining Brant with Ben-Michael, let alone Ben-Michael in combination with Kadambi.

Brant relates to a cache memory which is positioned between a host computer 10 and a disk storage system 25. Such uses of cache memory are well known. The advantage of a cache memory is that, with a typical memory access pattern that exhibits temporal and spatial locality, data that has been recently recovered from disk storage is likely to be accessed again in the near future. Such data is retained in the cache and subsequent accesses can be made directly to the cache without the need for an additional access to disk.

Since all data in a data stream must flow out of the switch in the same order in which it arrives, the simple FIFO memory architecture is well suited to the Ben-Michael application. One skilled in the art would see no need or advantage to the more complicated cache memory system of Brant *et al.* in a data streaming application like Ben-Michael. The cache system enables

-13-

unordered access to data made available on a probability basis, with the primary static disk storage always being available if there is a miss in the cache. In Ben-Michael, there is no static storage and there is no need for unordered access. Nor does Kadambi have static storage. Thus, there is no suggestion of combining the references.

Moreover, the three references are from completely different application areas. Kadambi describes memory management to store enroute packets for an Ethernet packet switch, Ben-Michael describes a credit FIFO for an ATM switch, and Brant describes a disk cache. The type of data being stored, the reference patterns for the data, and the operations being performed are completely different in each case.

A key aspect of the Avici invention is the application of caching (which is well known in other fields) to packet buffers in routers. This is a novel and non-obvious application of caching.

The Kadambi reference does not anticipate claim 1 of the Avici application because it does not operate its first set of buffers (CBP) as a cache, with information units being evicted to the second set of buffers (GBP). It does not do any eviction, in order of receipt or otherwise. Combining with Ben-Michael would not provide the missing element since Ben-Michael evicts, but only in order of receipt. Brant does disclose caching; however, the combination of Brant with Kadambi and Ben-Michael is non-obvious, and even if the references were combined, inventive input would be required to derive the claimed invention.

With respect to claims 3, 18, 33 and 41, which recited virtual channels, the Examiner has referred to the reference to ATM in Kadambi. However, as discussed in the last response, ATM switches make use of virtual circuits which are very distinct from the claimed virtual channels.

A virtual circuit in an ATM network is a path which is pre-established through multiple switch nodes for the flow of a data stream. A virtual channel, on the other hand, refers to buffer space which is required on a router to handle multiple data streams vying for a common destination port. In an ATM switch, data streams do not vie for an output port since the output

-14-

port has been pre-assigned to a virtual circuit through which the data stream is flowing. This distinction between virtual channels, which must arbitrate for an output port, and a virtual circuit is a key to the cache memory solution of the present invention as compared to the appropriate FIFO solution in the ATM switch environment. Virtual channels and the problem of allocating buffer space for the virtual channels do not exist in an ATM switch. There is no suggestion in any other cited reference of virtual channels. Accordingly, claims 3, 18, 33 and 41 should be allowed. Similarly, claims 7, 22 and 46-49 should be allowed.

With respect to claims 8-10, 23-25, 36, 37 and 44, which recited "flow control to stop the arrival of new information units while transferring information units between the first set of buffers," the Examiner has referred to the credit based flow control of Ben-Michael. However, in Ben-Michael it is the credits themselves which are being transferred to the FIFO, not the information units. One would not use flow control to stop flow of the flow control credits, and there is no suggestion of stopping the transfer of information units as the credits are transferred off-chip. Accordingly, it is respectfully submitted that those claims are not suggested by the cited references.

With respect to claims 11, 12, 26 and 27, which refer to miss status registers or an eviction buffer to store data prior to access or transfer to the second set of buffers, the Examiner has relied upon the FIFO bank 1 of Ben-Michael. First, from the description in Ben-Michael, the relationship between the FIFO banks 0 and 1 is not clear. It is clear, however, that credits are held in each bank primarily for transfer to the FIFO banks 2 and 3. Transfer to the local memory 720 is only in the event that space is not available in bank 2 or 3. Accordingly, the FIFO bank 1 must at best be considered part of the first set of rapidly accessible buffers and not a set of status registers or an eviction buffer whose function is to hold data being transferred to the off-chip memory. Accordingly, these claims are not suggested by the cited references.

With respect to claims 14-15, 30, 38-39 and 45, which refer to a "fabric router," the Examiner has referred to Kadambi. However, the Kadambi forwarding element is not a fabric router as that term has been used in the present application. As used in the present application, a

-15-

fabric router is a router that is within a larger switch or router. Such is not the case in Kadambi. Accordingly, those claims should be allowable.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By 
James M. Smith
Registration No. 28,043
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Dated: 7/22/04